

Greedy Algorithms for Optimal Distribution Approximation

Bernhard C. Geiger, Georg Böcherer

Institute for Communications Engineering, TU München, Germany

geiger@ieee.org, georg.boecherer@tum.de

Abstract—The approximation of a discrete probability distribution t by an M -type distribution p is considered. The approximation error is measured by the informational divergence $\mathbb{D}(t\|p)$, which is an appropriate measure, e.g., in the context of data compression. Properties of the optimal approximation are derived and bounds on the approximation error are presented, which are asymptotically tight. It is shown that M -type approximations that minimize either $\mathbb{D}(t\|p)$, or $\mathbb{D}(p\|t)$, or the variational distance $\|p - t\|_1$ can all be found by using specific instances of the same general greedy algorithm.

I. INTRODUCTION

In this work, we consider finite precision representations of probabilistic models. More precisely, if the original model, or *target distribution*, has n non-zero mass points and is given by $t := (t_1, \dots, t_n)$, we wish to approximate it by a distribution $p := (p_1, \dots, p_n)$, where for every i , $p_i = c_i/M$ for some non-negative integer $c_i \leq M$. The distribution p is called an *M-type distribution*, and the positive integer $M \geq n$ is the *precision* of the approximation. The problem is non-trivial, since computing the numerator c_i by rounding Mt_i to the nearest integer in general fails to yield a distribution.

M -type approximations have many practical applications, e.g., in political apportionments, M seats in a parliament need to be distributed to n parties according to the result of some vote t . This problem led, e.g., to the development of *multiplier methods* [1]. In communications engineering, example applications are finite precision implementations of probabilistic data compression [2], distribution matching [3], and finite-precision implementations of Bayesian networks [4], [5]. In all of these applications, the M -type approximation p should be close to the target distribution t in the sense of an appropriate error measure. Common choices for this approximation error are the variational distance and the informational divergences:

$$\|p - t\|_1 := \sum_{i=1}^n |p_i - t_i| \quad (1a)$$

$$\mathbb{D}(p\|t) := \sum_{i: p_i > 0} p_i \log \frac{p_i}{t_i} \quad (1b)$$

$$\mathbb{D}(t\|p) := \sum_{i: t_i > 0} t_i \log \frac{t_i}{p_i} \quad (1c)$$

where \log denotes the natural logarithm.

Variational distance and informational divergence (1b) have been considered by Reznik [6] and Böcherer [7], respectively, who presented algorithms for optimal M -type approximation

and developed bounds on the approximation error. In a recent manuscript [8], we extended the existing works on (1a) and (1b) to target distributions with infinite support ($n = \infty$) and refined the bounds from [6], [7].

In this work, we focus on the approximation error (1c). It is an appropriate cost function for data compression [9, Thm. 5.4.3] and seems appropriate for the approximation of parameters in Bayesian networks (see Sec. IV). Nevertheless, to the best of the authors' knowledge, the characterization of M -type approximations minimizing $\mathbb{D}(t\|p)$ has not received much attention in literature so far.

Our contributions are as follows. In Sec. II, we present an efficient greedy algorithm to find M -type distributions minimizing (1c). We then discuss in Sec. III the properties of the optimal M -type approximation and bound the approximation error (1c). Our bound incorporates a reverse Pinsker inequality recently suggested in [10, Thm. 7]. The algorithm we present is an instance of a greedy algorithm similar to *steepest ascent hill climbing* [11, Ch. 2.6]. As a byproduct, we unify this work with [6]–[8] by showing that also the algorithms optimal w.r.t. variational distance (1a) and informational divergence (1b) are instances of the same general greedy algorithm, see Sec. II.

II. GREEDY OPTIMIZATION

In this section, we define a class of problems that can be optimally solved by a greedy algorithm. Consider the following example:

Example 1. Suppose there are n queues with jobs, and you have to select M jobs minimizing the total time spent. A greedy algorithm suggests to select always the job with the shortest duration, among the jobs that are at the front of their queues. If the jobs in each queue are ordered by increasing duration, then this greedy algorithm is optimal.

We now make this precise: Let M be a positive integer, e.g., the number of jobs that have to be completed, and let $\delta_i: \mathbb{N} \rightarrow \mathbb{R}$, $i = 1, \dots, n$, be a set of functions, e.g., $\delta_i(k)$ is the duration of the k -th job in the i -th queue. Let furthermore $c_0 := (c_{1,0}, \dots, c_{n,0}) \in \mathbb{N}_0^n$ be a *pre-allocation*, representing a constraint that has to be fulfilled (e.g., in each queue at least one job has to be completed) or a chosen initialization. Then, the goal is to minimize

$$U(c) := \sum_{i=1}^n \sum_{k=c_{i,0}+1}^{c_i} \delta_i(k_i) \quad (2)$$

Algorithm 1. Greedy Algorithm

Initialize $k_i = c_{i,0}$, $i = 1, \dots, n$.
repeat $M - \|\mathbf{c}_0\|_1$ times
 Compute $\delta_i(k_i + 1)$, $i = 1, \dots, n$.
 Compute $j = \min \arg \min_i \delta_i(k_i + 1)$.
 // (choose one minimal element)
 Update $k_j \leftarrow k_j + 1$.
end repeat
Return $\mathbf{c} = (k_1, \dots, k_n)$.

i.e., to find a *final allocation* $\mathbf{c} := (c_1, \dots, c_n)$ satisfying, for all i , $c_i \geq c_{i,0}$ and $\|\mathbf{c}\|_1 = M$. A greedy method to obtain such a final allocation is presented in Algorithm 1. We show in Appendix A that this algorithm is optimal if the functions δ_i satisfy certain conditions:

Proposition 1. *If the functions $\delta_i(k)$ are non-decreasing in k , Algorithm 1 achieves a global minimum $U(\mathbf{c})$ for a given pre-allocation \mathbf{c}_0 and a given M .*

Remark 1. *The minimum of $U(\mathbf{c})$ may not be unique.*

Remark 2. *If a function $f_i: \mathbb{R} \rightarrow \mathbb{R}$ is convex, the difference $\delta_i(k) = f_i(k) - f_i(k-1)$ is non-decreasing in k . Hence, Algorithm 1 also minimizes*

$$U(\mathbf{c}) = \sum_{i=1}^n f_i(c_i). \quad (3)$$

Remark 2 connects Algorithm 1 to steepest ascent hill climbing [11, Ch. 2.6] with fixed step size and a constrained number of M steps. Hill climbing is optimal for convex problems, suggesting an interesting connection with Proposition 1.

We now show that instances of Algorithm 1 can find M -type approximations \mathbf{p} minimizing each of the cost functions in (1). Noting that $p_i = c_i/M$ for some non-negative integer c_i , we can rewrite the cost functions as follows:

$$\|\mathbf{p} - \mathbf{t}\|_1 = \frac{1}{M} \sum_{i=1}^n |c_i - Mt_i| \quad (4a)$$

$$\mathbb{D}(\mathbf{p}||\mathbf{t}) = \frac{1}{M} \left(\sum_{i: c_i > 0} c_i \log \frac{c_i}{t_i} \right) - \log M \quad (4b)$$

$$\mathbb{D}(\mathbf{t}||\mathbf{p}) = \log M - H(\mathbf{t}) - \sum_{i: t_i > 0} t_i \log c_i. \quad (4c)$$

Ignoring constant terms, these cost functions are all instances of Remark 2 for convex functions $f_i: \mathbb{R} \rightarrow \mathbb{R}$ (see Table I). Hence, the three different M -type approximation problems set up by (1) can all be solved by instances of Algorithm 1, for a trivial pre-allocation $\mathbf{c}_0 = \mathbf{0}$ and after taking M steps: The final allocation \mathbf{c} simply defines the M -type approximation by $p_i = c_i/M$. For variational distance optimal approximation, we showed in [8, Lem. 3] that every optimal M -type approximation satisfies $p_i \geq \lfloor Mt_i \rfloor / M$, hence one may speed up the algorithm by pre-allocating $c_{i,0} = \lfloor Mt_i \rfloor$. We furthermore show in Lemma 1 below that the support of the optimal M -type approximation in terms of (1c) equals the support of \mathbf{t} (if M is large enough). Assuming that \mathbf{t} is positive, one can

TABLE I
INSTANCES OF ALGORITHM 1 OPTIMIZING (1).

Cost	$f_i(x)$	$\delta_i(k)$	$c_{i,0}$	Refs
$\ \mathbf{p} - \mathbf{t}\ _1$	$ x - Mt_i $	$ k - Mt_i $ $- k-1 - Mt_i $	$\lfloor Mt_i \rfloor$	[6], [8]
$\mathbb{D}(\mathbf{p} \mathbf{t})$	$x \log(x/t_i)$	$k \log \frac{k}{k-1}$ $+ \log(k-1)$ $- \log t_i$	0	[7], [8]
$\mathbb{D}(\mathbf{t} \mathbf{p})$	$-t_i \log x$	$t_i \log((k-1)/k)$	$\lceil t_i \rceil$	This work

pre-allocate the algorithm with $c_{i,0} = 1$. We summarize these instantiations of Algorithm 1 in Table I.

This list of instances of Algorithm 1 minimizing information-theoretic or probabilistic cost functions can be extended. For example, the χ^2 -divergences $\chi^2(\mathbf{t}||\mathbf{p})$ and $\chi^2(\mathbf{p}||\mathbf{t})$ can also be minimized, since the functions inside the respective sums are convex. However, Rényi divergences of orders $\alpha \neq 1$ cannot be minimized by applying Algorithm 1.

III. M -TYPE APPROXIMATION MINIMIZING $\mathbb{D}(\mathbf{t}||\mathbf{p})$: PROPERTIES AND BOUND

As shown in the previous section, Algorithm 1 presents a minimizer of the problem $\min_{\mathbf{p}} \mathbb{D}(\mathbf{t}||\mathbf{p})$ if instantiated according to Table I. Let us call this minimizer \mathbf{t}^a . Recall that \mathbf{t} is positive and that $M \geq n$. The support of \mathbf{t}^a must be at least as large as the support of \mathbf{t} , since otherwise $\mathbb{D}(\mathbf{t}||\mathbf{t}^a) = \infty$. Note further that the costs $\delta_i(k)$ are negative if $t_i > 0$ and zero if $t_i = 0$; hence, if $t_i = 0$, the index i cannot be chosen by Algorithm 1, thus also $t_i^a = 0$. This proves

Lemma 1. *If $M \geq n$, the supports of \mathbf{t} and \mathbf{t}^a coincide, i.e., $t_i = 0 \Leftrightarrow t_i^a = 0$.*

The assumption that \mathbf{t} is positive and that $M \geq n$ hence comes without loss of generality. In contrast, neither variational distance nor informational divergence (1b) require $M \geq n$: As we show in [8], the M -type approximation problem remains interesting even if $M < n$.

Lemma 1 explains why the optimal M -type approximation does not necessarily result in a “small” approximation error:

Example 2. Let $\mathbf{t} = (1 - \varepsilon, \frac{\varepsilon}{n-1}, \dots, \frac{\varepsilon}{n-1})$ and $M = n$, hence by Lemma 1, $\mathbf{t}^a = \frac{1}{n}(1, 1, \dots, 1)$. It follows that $\mathbb{D}(\mathbf{t}||\mathbf{t}^a) = \log n - H(\mathbf{t})$, which can be made arbitrarily close to $\log n$ by choosing a small positive ε .

In Table I we made use of [8, Lem. 3], which says that every \mathbf{p} minimizing the variational distance $\|\mathbf{p} - \mathbf{t}\|_1$ satisfies $p_i \geq \lfloor Mt_i \rfloor / M$, to speed up the corresponding instance of Algorithm 1 by proper pre-allocation. Initialization by rounding is not possible when minimizing $\mathbb{D}(\mathbf{t}||\mathbf{p})$, as shown in the following two examples:

Example 3. Let $\mathbf{t} = (17/20, 3/40, 3/40)$ and $M = 20$. The optimal M -type approximation is $\mathbf{p} = (8/10, 1/10, 1/10)$, hence $p_1 < \lfloor Mt_1 \rfloor / M$. Initialization via rounding off fails.

Example 4. Let $\mathbf{t} = (0.719, 0.145, 0.088, 0.048)$ and $M = 50$. The optimal M -type approximation is $\mathbf{p} =$

(0.74, 0.14, 0.08, 0.04), hence $p_1 > \lceil Mt_1 \rceil / M$. Initialization via rounding up fails.

To show that informational divergence vanishes for $M \rightarrow \infty$, assume that $M > 1/t_i$ for all i . Hence, the variational distance optimal approximation \mathbf{t}^{vd} has the same support as \mathbf{t} , which ensures that $\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{vd}}) < \infty$. By similar arguments as in the proof of [8, Prop. 4.1)], we obtain

$$\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{a}}) \leq \mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{vd}}) \leq \log \left(1 + \frac{n}{2M} \right) \xrightarrow{M \rightarrow \infty} 0. \quad (5)$$

We now develop an upper bound on $\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{a}})$ that holds for every M . To this end, we first approximate \mathbf{t} by a distribution \mathbf{t}^* in $\mathcal{P}_M := \{\mathbf{p}: \forall i: p_i \geq 1/M, \|\mathbf{p}\|_1 = 1\}$ that minimizes $\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^*)$. If \mathbf{t}^* is unique, then it is called the *reverse I-projection* [12, Sec. I.A] of \mathbf{t} onto \mathcal{P}_M . Since $\mathbf{t}^* \in \mathcal{P}_M$, its variational distance optimal approximation \mathbf{t}^{vd} has the same support as \mathbf{t} , which allows us to bound $\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{a}})$ by $\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{vd}})$.

Lemma 2. Let $\mathbf{t}^* \in \mathcal{P}_M$ minimize $\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^*)$. Then,

$$t_i^* := \frac{t_i}{\nu(M)} + \left(\frac{1}{M} - \frac{t_i}{\nu(M)} \right)^+ \quad (6)$$

where $\nu(M)$ is such that $\|\mathbf{t}^*\|_1 = 1$, and where $(x)^+ := \max\{0, x\}$.

Proof: See Appendix B. ■

Let $\mathcal{K} := \{i : t_i < \nu(M)/M\}$, $k := |\mathcal{K}|$, and $T_{\mathcal{K}} := \sum_{i \in \mathcal{K}} t_i$. The parameter ν must scale the mass $(1 - T_{\mathcal{K}})$ such that it equals $(M - k)/M$, i.e., we have

$$\nu = \frac{1 - T_{\mathcal{K}}}{1 - \frac{k}{M}}. \quad (7)$$

If, for all i , $t_i > 1/M$, then $\mathbf{t} \in \mathcal{P}_M$, hence $\mathbf{t}^* = \mathbf{t}$ is feasible and $\nu(M) = 1$. One can show that $\nu(M)$ decreases with M .

Proposition 2 (Approximation Bounds).

$$\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{a}}) \leq \log \nu(M) + \frac{\log(2)}{2} \left(1 - \nu(M) \left(1 - \frac{n}{M} \right) \right) \quad (8)$$

Proof: See Appendix C. ■

The first term on the right-hand side of (8) accounts for the error caused by first approximating \mathbf{t} by \mathbf{t}^* (in the sense of Lemma 2). The second term then accounts for the additional error caused by the M -type approximation of \mathbf{t}^* . The bound incorporates the reverse Pinsker inequality [10, Thm. 7], see Appendix C. If $M > t_i$ for every i , hence $\mathbf{t} \in \mathcal{P}_M$, then $\nu(M) = 1$ and only the second term remains. For large M , (8) yields better results than (5), justifying this bound on the whole range of M . We illustrate the bounds for an example in Fig. 1.

IV. OUTLOOK

A possible application for our algorithms is the M -type approximation of Markov models, i.e., approximating the transition matrix \mathbf{T} of an n -state, irreducible Markov chain with invariant distribution vectors $\boldsymbol{\mu}$ by a transition matrix \mathbf{P} containing only M -type probabilities. Generalizing (1c), the approximation error can be measured by the informational divergence rate [13]

$$\overline{\mathbb{D}}(\mathbf{T} \parallel \mathbf{P}) := \sum_{i,j=1}^n \mu_i T_{ij} \log \frac{T_{ij}}{P_{ij}} = \sum_{i=1}^n \mu_i \mathbb{D}(\mathbf{t}_i \parallel \mathbf{p}_i) \quad (9)$$

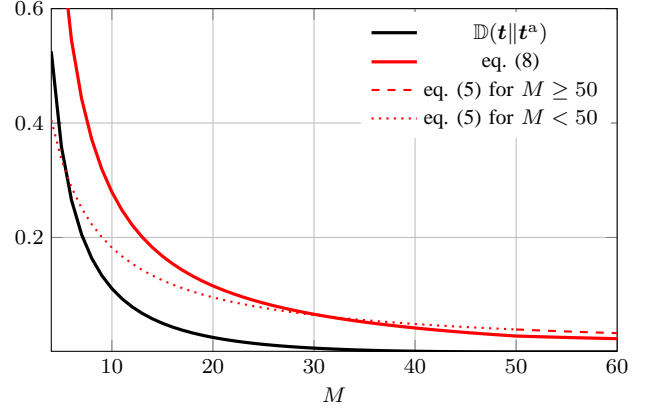


Fig. 1. Evaluating the bounds (8) and (5) for $\mathbf{t} = (0.48, 0.48, 0.02, 0.02)$. Note that (5) is a valid bound only for $M \geq 50$, i.e., where the curve is dashed.

The optimal M -type approximation is found by applying the instance of Algorithm 1 to each row separately, and Lemma 1 ensures that the transition graph of \mathbf{P} equals that of \mathbf{T} , i.e., the approximating Markov chain is irreducible. Future work shall extend this analysis to hidden Markov models and should investigate the performance of these algorithms in practical scenarios, e.g., in speech processing.

Another possible application is the approximation of Bayesian network parameters. The authors of [4] approximated the true parameters using a stationary multiplier method from [14]. Since rounding probabilities to zero led to bad classification performance, they replaced zeros in the approximating distribution afterwards by small values; thus, they also approximated true zero probabilities by positive ones. We believe that these problems can be removed by instantiating Algorithm 1 for cost (1c). This automatically prevents approximating non-zero probabilities with zeros and vice-versa, see Lemma 1.

Recent work suggested rounding *log-probabilities*, i.e., to approximate $\log t_i$ by $\log p_i = -c_i/M$ for a non-negative integer c_i [5]. Finding an optimal approximation that corresponds to a true distribution is equivalent to solving

$$\begin{aligned} \min \quad & d(\mathbf{t}, \mathbf{p}) \\ \text{s.t.} \quad & \|e^{-c}\|_{1/M} = 1 \end{aligned}$$

where $d(\cdot, \cdot)$ denotes any of the considered cost functions (1). If $M = 1$ and $d(\mathbf{t}, \mathbf{p}) = \mathbb{D}(\mathbf{t} \parallel \mathbf{p})$ using the binary logarithm, the constraint translates to the requirement that \mathbf{t} is approximated by a complete binary tree. Then, the optimal approximation is the Huffman code for \mathbf{t} .

ACKNOWLEDGEMENTS

The work of Bernhard C. Geiger was partially funded by the Erwin Schrödinger Fellowship J 3765 of the Austrian Science Fund. The work of Georg Böcherer was partly supported by the German Ministry of Education and Research in the framework of an Alexander von Humboldt Professorship.

APPENDIX

A. Proof of Proposition 1

Since a pre-allocation only fixes a lower bound for $U(\mathbf{c})$, w.l.o.g. we assume that $\mathbf{c}_0 = \mathbf{0}$ and thus $\mathbf{c} \in \mathbb{N}_0^n$ with $\|\mathbf{c}\|_1 = M$. Consider the set $\mathcal{D} := \{\delta_i(k_i) : k_i \in \mathbb{N}, i = 1, \dots, n\}$ and assume that the (not necessarily unique) set \mathcal{D}_M consists of M smallest values in \mathcal{D} , i.e., $|\mathcal{D}_M| = M$ and

$$\forall d \in \mathcal{D}_M, d' \in \mathcal{D} \setminus \mathcal{D}_M: d \leq d'. \quad (10)$$

Clearly, $U(\mathbf{c})$ cannot be smaller than the sum over all elements in \mathcal{D}_M . Since the δ_i are non-decreasing, there exists at least one final allocation $\tilde{\mathbf{c}}$ that takes successively the first c_i values from each queue i , i.e., $\mathcal{D}_M = \{\delta_1(1), \dots, \delta_1(c_1), \dots, \delta_n(1), \dots, \delta_n(c_n)\}$ satisfies (10). This shows that the lower bound induced by (10) can actually be achieved.

We prove the optimality of Algorithm 1 by contradiction: Assume that Algorithm 1 finishes with a final allocation $\tilde{\mathbf{c}}$ such that $U(\tilde{\mathbf{c}})$ is strictly larger than the (unique) sum over all elements in (non-unique) \mathcal{D}_M . Hence, $\tilde{\mathbf{c}}$ must exchange at least one of the elements in \mathcal{D}_M for an element that is strictly larger. Thus, by the properties of the functions δ_i and Algorithm 1, there must be indices ℓ and m such that $\tilde{c}_\ell > c_\ell$, $\tilde{c}_m < c_m$, and $\delta_\ell(\tilde{c}_\ell) \geq \delta_\ell(c_\ell + 1) > \delta_m(c_m) \geq \delta_m(\tilde{c}_m)$. At each iteration of the algorithm, the current allocation at index m satisfies $k_m \leq \tilde{c}_m < c_m$. Since $\delta_m(c_m) < \delta_\ell(c_\ell + 1)$, $\delta_\ell(c_\ell + 1)$ can never be a minimal element, and hence is not chosen by Algorithm 1. This contradicts the assumption that Algorithm 1 finishes with a $\tilde{\mathbf{c}}$ such that $U(\tilde{\mathbf{c}})$ is strictly larger than the sum of \mathcal{D} 's M smallest values. \square

B. Proof of Lemma 2

The problem finding a $\mathbf{t}^* \in \mathcal{P}_M$ minimizing $\mathbb{D}(\mathbf{t} \|\mathbf{t}^*)$ is equivalent to finding an optimal point of the problem:

$$\underset{\mathbf{p} \in \mathbb{R}_{>0}^n}{\text{minimize}} \quad - \sum_{i=1}^n t_i \log p_i \quad (11a)$$

$$\text{subject to} \quad \frac{1}{M} - p_i \leq 0, \quad i = 1, 2, \dots, n \quad (11b)$$

$$-1 + \sum_{i=1}^n p_i = 0 \quad (11c)$$

The Lagrangian of the problem is

$$\begin{aligned} L(\mathbf{p}, \boldsymbol{\lambda}, \nu) = & - \sum_{i=1}^n t_i \log p_i \\ & + \sum_{i=1}^n \lambda_i \left(\frac{1}{M} - p_i \right) + \nu \left(-1 + \sum_{i=1}^n p_i \right). \end{aligned} \quad (12)$$

By the KKT conditions [15, Ch. 5.5.3], a feasible point \mathbf{t}^* is optimal if, for every $i = 1, \dots, n$,

$$\lambda_i \geq 0 \quad (13a)$$

$$\lambda_i \left(\frac{1}{M} - t_i^* \right) = 0 \quad (13b)$$

$$\frac{\partial}{\partial p_i} L(\mathbf{p}, \boldsymbol{\lambda}, \nu)|_{\mathbf{p}=\mathbf{t}^*} = -\frac{t_i}{t_i^*} - \lambda_i + \nu = 0 \quad (13c)$$

By (11b), we have $t_i^* \geq 1/M$. If $t_i^* > 1/M$, then $\lambda_i = 0$ by (13b) and $t_i^* = t_i/\nu$ by (13c). Thus

$$t_i^* = \frac{t_i}{\nu} + \left(\frac{1}{M} - \frac{t_i}{\nu} \right)^+ \quad (14)$$

where ν is such that $\sum_{i=1}^n t_i^* = 1$. \square

C. Proof of Proposition 2

Reverse I -projections admit a Pythagorean inequality [12, Thm. 1]. In other words, if \mathbf{p} is a distribution, \mathbf{p}^* its reverse I -projection onto a set \mathcal{S} , and \mathbf{q} any distribution in \mathcal{S} , then

$$\mathbb{D}(\mathbf{p} \|\mathbf{q}) \geq \mathbb{D}(\mathbf{p} \|\mathbf{p}^*) + \mathbb{D}(\mathbf{p}^* \|\mathbf{q}). \quad (15)$$

For the present scenario, we can show an even stronger result:

Lemma 3. *Let \mathbf{t} be the target distribution, let \mathbf{t}^* be as in Lemma 2, and let \mathbf{t}^{vd} be the variational distance optimal M -type approximation of \mathbf{t}^* . Then,*

$$\mathbb{D}(\mathbf{t} \|\mathbf{t}^{\text{vd}}) = \mathbb{D}(\mathbf{t} \|\mathbf{t}^*) + \nu \mathbb{D}(\mathbf{t}^* \|\mathbf{t}^{\text{vd}}). \quad (16)$$

Proof:

$$\mathbb{D}(\mathbf{t} \|\mathbf{t}^{\text{vd}}) = \sum_{i=1}^n t_i \log \frac{t_i}{t_i^{\text{vd}}} \quad (17)$$

$$= \sum_{i=1}^n t_i \log \frac{t_i t_i^*}{t_i^{\text{vd}} t_i^*} \quad (18)$$

$$= \sum_{i=1}^n t_i \log \frac{t_i}{t_i^*} + \sum_{i=1}^n t_i \log \frac{t_i^*}{t_i^{\text{vd}}} \quad (19)$$

$$\stackrel{(a)}{=} \mathbb{D}(\mathbf{t} \|\mathbf{t}^*) + \nu \sum_{i \notin \mathcal{K}} \frac{t_i}{\nu} \log \frac{t_i^*}{t_i^{\text{vd}}} \quad (20)$$

$$\stackrel{(b)}{=} \mathbb{D}(\mathbf{t} \|\mathbf{t}^*) + \nu \mathbb{D}(\mathbf{t}^* \|\mathbf{t}^{\text{vd}}) \quad (21)$$

where (a) follows because for $i \in \mathcal{K}$, $t_i^* = t_i^{\text{vd}} = 1/M$ and (b) is because for $i \notin \mathcal{K}$, $t_i^* = t_i/\nu$. \blacksquare

We now bound the summands in Lemma 3.

Lemma 4. *In the setting of Lemma 3,*

$$\mathbb{D}(\mathbf{t}^* \|\mathbf{t}^{\text{vd}}) \leq \log(2) \|\mathbf{t}^* - \mathbf{t}^{\text{vd}}\|_1. \quad (22)$$

Proof: We first employ a reverse Pinsker inequality from [10, Thm. 7], stating that

$$\mathbb{D}(\mathbf{t}^* \|\mathbf{t}^{\text{vd}}) \leq \frac{1}{2} \frac{r \log r}{r-1} \|\mathbf{t}^* - \mathbf{t}^{\text{vd}}\|_1 \quad (23)$$

where $r := \sup_{i: t_i^* > 0} \frac{t_i^*}{t_i^{\text{vd}}}$. Furthermore, since for variational distance optimal approximations we always have $|t_i^* - t_i^{\text{vd}}| < 1/M$ [8, Lem. 3], we can bound

$$r < \frac{t_i^{\text{vd}} + \frac{1}{M}}{t_i^{\text{vd}}} \leq 2 \quad (24)$$

since $t_i^{\text{vd}} \geq \lfloor M t_i^* \rfloor / M \geq 1/M$. \blacksquare

Lemma 5. *In the setting of Lemma 3,*

$$\mathbb{D}(\mathbf{t} \|\mathbf{t}^*) \leq \log \nu. \quad (25)$$

Proof:

$$\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^*) = \sum_{i=1}^n t_i \log \frac{t_i}{t_i^*} \quad (26)$$

$$= \sum_{i \notin \mathcal{K}} t_i \log \frac{\nu t_i}{t_i} + \sum_{i \in \mathcal{K}} t_i \log M t_i \quad (27)$$

$$\stackrel{(a)}{\leq} (1 - T_{\mathcal{K}}) \log \nu + \sum_{i \in \mathcal{K}} t_i \log \nu \quad (28)$$

$$= \log \nu \quad (29)$$

where (a) is because for $i \in \mathcal{K}$, $M t_i \leq \nu$. ■

To bound $\|\mathbf{t}^* - \mathbf{t}^{\text{vd}}\|_1$, we present

Lemma 6. *Let \mathbf{p}^* be a sub-probability distribution with $m \leq M$ masses and total weight $1 - T$, and let $\mathbf{p}^{\text{vd}*}$ be the variational distance optimal M -type approximation using $J \leq M$ masses. Then,*

$$\|\mathbf{p}^* - \mathbf{p}^{\text{vd}*}\|_1 \leq \frac{m}{2M} + \frac{(M - MT - J)^2}{2mM}. \quad (30)$$

Note that for $J = M$ we recover [8, Lemma 4].

Proof: Assume that either $\forall i: p_i^* \geq p_i^{\text{vd}*}$ or $\forall i: p_i^* \leq p_i^{\text{vd}*}$. Note that this is possible since \mathbf{p}^* and $\mathbf{p}^{\text{vd}*}$ are sub-probability distributions. Then, $\|\mathbf{p}^* - \mathbf{p}^{\text{vd}*}\|_1 = |1 - T - J/M|$ which satisfies this bound. This can be seen by rearranging (30) such that J only appears on the left-hand side; the maximizing J (not necessarily integer) then satisfies (30) with equality.

We thus remain to treat the case where after rounding off all indices, $1 \leq L \leq M - 1$ masses remain and we have

$$\sum_{i=1}^m p_i^* - \frac{\lfloor M p_i^* \rfloor}{M} =: \sum_{i=1}^m e_i = 1 - T - \frac{J - L}{M} =: g(L). \quad (31)$$

The variational distance is minimized by distributing the L masses to L indices $i \in \mathcal{L}$ with the largest errors e_i , hence

$$\|\mathbf{p}^* - \mathbf{p}^{\text{vd}*}\|_1 = \sum_{i \in \mathcal{L}} \left(\frac{1}{M} - e_i \right) + \sum_{i \notin \mathcal{L}} e_i \quad (32)$$

$$\stackrel{(a)}{\leq} \frac{L}{M} - \frac{L}{n} g(L) + \frac{n - L}{n} g(L) \quad (33)$$

where (a) follows because for $i \in \mathcal{L}, j \notin \mathcal{L}$, $e_i \geq e_j$. This is maximized for $L = \frac{n - (M - MT - J)}{2}$ (not necessarily integer), which after inserting yields the upper bound. ■

Proof of Bound in Proposition 2: We start by bounding the informational divergence $\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^a)$ by the informational divergence between \mathbf{t} and the variational distance optimal approximation \mathbf{t}^{vd} of its reverse I -projection \mathbf{t}^* onto \mathcal{P}_M :

$$\mathbb{D}(\mathbf{t} \parallel \mathbf{t}^a) \leq \mathbb{D}(\mathbf{t} \parallel \mathbf{t}^{\text{vd}}) \quad (34)$$

$$\stackrel{(a)}{=} \mathbb{D}(\mathbf{t} \parallel \mathbf{t}^*) + \nu \mathbb{D}(\mathbf{t}^* \parallel \mathbf{t}^{\text{vd}}) \quad (35)$$

$$\stackrel{(b)}{\leq} \log \nu + \nu \log(2) \|\mathbf{t}^* - \mathbf{t}^{\text{vd}}\|_1 \quad (36)$$

$$\stackrel{(c)}{\leq} \log \nu + \nu \log(2) \frac{n - k}{2M} \quad (37)$$

$$\stackrel{(d)}{\leq} \log \nu + \nu \log(2) \frac{n - M + \frac{M}{\nu}}{2M} \quad (38)$$

$$= \log \nu + \frac{\log(2)}{2} \left(1 - \nu \left(1 - \frac{n}{M} \right) \right) \quad (39)$$

where

- (a) is due to Lemma 3,
- (b) is due to Lemmas 4 and 5,
- (c) is due to Lemma 6 with $m = n - k$, $1 - T = 1 - k/M$, and $J = M - k$, and
- (d) follows by bounding k from below via (7)

$$k = \frac{M}{\nu} (\nu - 1 + T_{\mathcal{K}}) \geq \frac{M}{\nu} (\nu - 1) = M - \frac{M}{\nu}. \quad (40)$$

■

REFERENCES

- [1] G. Dörflleitner and T. Klein, “Rounding with multiplier methods: An efficient algorithm and applications in statistics,” *Statistical Papers*, vol. 40, pp. 143–157, 1999.
- [2] J. Rissanen and G. G. Langdon, “Arithmetic coding,” *IBM J. Res. Develop.*, vol. 23, no. 2, pp. 149–162, 1979.
- [3] P. Schulte and G. Böcherer, “Constant composition distribution matching,” *IEEE Trans. Inf. Theory*, vol. 62, no. 1, pp. 430–434, Jan. 2016.
- [4] M. J. Druzdżel and A. Oniško, “Are Bayesian networks sensitive to precision of their parameters?” in *Proc. Int. Conf. Intelligent Information Systems*, Zakopane, Jun. 2008, pp. 35–44.
- [5] S. Tschatschek and F. Pernkopf, “On Bayesian network classifiers with reduced precision parameters,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 4, pp. 774–785, Apr. 2015.
- [6] Y. Reznik, “An algorithm for quantization of discrete probability distributions,” in *Data Compression Conference (DCC)*, Snowbird, UT, Mar. 2011, pp. 333–342.
- [7] G. Böcherer, “Optimal non-uniform mapping for probabilistic shaping,” in *Proc. Int. ITG Conf. Systems, Communication and Coding (SCC)*, Munich, Jan. 2013, pp. 1–6.
- [8] G. Böcherer and B. C. Geiger, “Optimal quantization for distribution synthesis,” submitted to *IEEE Trans. Inf. Theory*; preprint available: arXiv:1307.6843 [cs.IT].
- [9] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Hoboken, NJ: Wiley Interscience, 2006.
- [10] S. Verdú, “Total variation distance and the distribution of relative information,” in *Proc. Inf. Theory and Applicat. Workshop (ITA)*, San Diego, CA, Feb. 2014, pp. 499–501.
- [11] Z. Michalewicz and D. B. Fogel, *How to Solve It: Modern Heuristics*, 2nd ed. Berlin: Springer, 2004.
- [12] I. Csiszár and M. František, “Information projections revisited,” *IEEE Trans. Inf. Theory*, vol. 49, no. 6, pp. 1474–1490, Jun. 2003.
- [13] Z. Rached, F. Alajaji, and L. L. Campbell, “The Kullback-Leibler divergence rate between Markov sources,” *IEEE Trans. Inf. Theory*, vol. 50, no. 5, pp. 917–921, May 2004.
- [14] L. Heinrich, F. Pukelsheim, and U. Schwingenschlögl, “On stationary multiplier methods for the rounding of probabilities and the limiting law of the Sainte-Laguë divergence,” *Statistics & Decisions*, vol. 23, pp. 117–129, 2005.
- [15] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge et al.: Cambridge University Press, 2004.